

WHAT IS CLAIMED IS:

1. A method of managing the availability to do work of a plurality of cooperating computers connected in a network, said method comprising:
 - (a) identifying a set of specific ones of said plurality of cooperating computers as available resources for the performance of said work;
 - (b) receiving performance related values of said plurality of cooperating computers;
 - (c) deriving performance related metrics from said performance related values; and
 - (d) changing said set of specific ones of said plurality of cooperating computers based on said performance related metrics.
2. The method of claim 1, wherein at least one of said plurality of cooperating computers is heterogeneous with respect to the other cooperating computers of said plurality of cooperating computers.
3. The method of claim 1, wherein step (d) adds additional ones of said plurality of cooperating computers to said set or deletes one or more of said specific ones of said plurality of cooperating computers from said set.
4. The method of claim 1, wherein step (d) changes said set independently of any architecture or operating system specific metrics of said plurality of cooperating computers.
5. The method of claim 1, wherein step (d) changes said set independently of any workload specific metrics of said plurality of cooperating computers.

6. The method of claim 1, wherein said performance values are selected from the group consisting of: response times and queue delays.
7. The method of claim 1, wherein a cluster of said plurality of cooperating computers is connected to a node contained in said network, further comprising (e) requesting a manager of said cluster to accept additional work or to give up pending work based on said performance related metrics.
8. The method of claim 7, further comprising (f) requesting said manager of said cluster to start more work or to run more pieces of an application on one or more of the cooperating computers of said cluster.
9. The method of claim 1, wherein step (a) identifies said set at a first time based on said performance related metrics, and wherein step (d) changes said set at a second later time.
10. The method of claim 1, wherein step (d) is performed only when a new value has been received or a request has been made to view the data.
11. The method of claim 1, wherein step (d) forms said performance metrics as an aggregation of said values.
12. The method of claim 11, wherein step (d) forms said performance metrics for each of said plurality of said cooperating computers.
13. The method of claim 12, wherein step (d) is performed only when a new one of said values is received or a request to view the performance metric is received.

14. The method of claim 12, wherein step (b) receives said performance related values over a series of time intervals, and wherein step (c) derives said performance metrics for n periods, of which the performance metrics of the nth period thereof includes an aggregate of the performance metrics for a current interval plus n-1 of the preceding intervals.
15. The method of claim 14, wherein the performance metrics of the nth period of a preceding interval are discarded during a current interval.
16. The method of claim 15, wherein said performance metrics for each of said periods include only a number and average of values received.
17. The method of claim 1, wherein each of said performance metrics includes only a number and an average of values received.
18. The method of claim 14, wherein step (c) forms said performance metrics as a data structure having n rows that contain the performance metrics of said n periods, respectively, wherein the performance metrics of the nth row of a preceding interval are discarded during a current interval, and wherein said nth row of the preceding interval is used as a first row in the current interval and the remaining ones of said n rows are shifted down one row position.
19. A computer having a CPU and a memory comprising:

policy program means for causing said CPU to manage the availability to do work of a plurality of cooperating computers that are connected in a network, said policy program means comprising:

first means for performing a first operation that identifies a set of specific ones of said plurality of cooperating computers as available resources for the performance of work;

second means for performing a second operation that receives performance related values of said plurality of cooperating computers;

third means for performing a third operation that derives performance related metrics from said performance related values; and

fourth means for performing a fourth operation that changes said set of specific ones of said plurality of cooperating computers based on said performance related metrics.

20. The computer of claim 19, wherein at least one of said plurality of cooperating computers is heterogeneous with respect to the other cooperating computers of said plurality of cooperating computers.

21. The computer of claim 19, wherein said fourth operation adds additional ones of said plurality of cooperating computers to said set or deletes one or more of said specific ones of said plurality of cooperating computers from said set.

22. The computer of claim 19, wherein said fourth operation changes said set independently of any architecture or operating system specific metrics of said plurality of cooperating computers.

23. The computer of claim 19, wherein said fourth operation changes said set independently of any workload specific metrics of said plurality of cooperating computers.

24. The computer of claim 19, wherein said performance values are selected from the group consisting of: response times and queue delays.

25. The computer of claim 19, wherein a cluster of said plurality of cooperating computers is connected to a node contained in said network, further comprising a fifth means for performing a fifth operation that requests a manager of said cluster to accept additional work or to give up pending work based on said performance related metrics.

26. The computer of claim 25, further comprising sixth means for performing a sixth operation that requests said manager of said cluster to start more work or to run more pieces of an application on one or more of the cooperating computers of said cluster.

27. The computer of claim 19, wherein said first operation identifies said set at a first time based on said performance related metrics, and wherein said fourth changes said set at a second later time.

28. The computer of claim 19, wherein said fourth operation is performed only when a new value has been received or a request has been made to view the data.

29. The computer of claim 19, wherein said fourth operation forms said performance metrics as an aggregation of said values.

30. The computer of claim 29, wherein said fourth operation forms said performance metrics for each of said plurality of said cooperating computers.

31. The computer of claim 30, wherein said fourth operation is performed only when a new one of said values is received or a request to view the performance metric is received.

32. The computer of claim 30, wherein said second operation receives said performance related values over a series of time intervals, and wherein said third

operation derives said performance metrics for n periods, of which the performance metrics of the n th period thereof includes an aggregate of the performance metrics for a current interval plus $n-1$ of the preceding intervals.

33. The computer of claim 32, wherein the performance metrics of the n th period of a preceding interval are discarded during a current interval.

34. The computer of claim 33, wherein said performance metrics for each of said periods include only a number and average of values received.

35. The computer of claim 19, wherein each of said performance metrics includes only a number and an average of values received.

36. The computer of claim 32, wherein said third operation forms said performance metrics as a data structure having n rows that contain the performance metrics of said n periods, respectively, wherein the performance metrics of the n th row of a preceding interval are discarded during a current interval, and wherein said n th row of the preceding interval is used as a first row in the current interval and the remaining ones of said n rows are shifted down one row position.

37. A memory medium for controlling a computer to manage the availability to do work of a plurality of cooperating computers connected in a network, said memory medium comprising:

first means for controlling said computer to perform a first operation that identifies a set of specific ones of said plurality of cooperating computers as available resources for the performance of work;

second means for controlling said computer to perform a second operation that receives performance related values of said plurality of cooperating computers;

third means for controlling said computer to perform a third operation that derives performance related metrics from said performance related values; and

fourth means for controlling said computer to perform a fourth operation that changes said set of specific ones of said plurality of cooperating computers based on said performance related metrics.

38. The memory medium of claim 37, wherein at least one of said plurality of cooperating computers is heterogeneous with respect to the other cooperating computers of said plurality of cooperating computers.

39. The memory medium of claim 37, wherein said fourth operation adds additional ones of said plurality of cooperating computers to said set or deletes one or more of said specific ones of said plurality of cooperating computers from said set.

40. The memory medium of claim 37, wherein said fourth operation changes said set independently of any architecture or operating system specific metrics of said plurality of cooperating computers.

41. The memory medium of claim 37, wherein said fourth operation changes said set independently of any workload specific metrics of said plurality of cooperating computers.

42. The memory medium of claim 37, wherein said performance values are selected from the group consisting of: response times and queue delays.

43. The memory medium of claim 37, wherein a cluster of said plurality of cooperating computers is connected to a node contained in said network, further comprising fifth means for performing a fifth operation that requests a manager of said cluster to accept additional work or to give up pending work based on said performance related metrics.

44. The memory medium of claim 43, further comprising sixth means for performing a sixth operation that requests said manager of said cluster to start more work or to run more pieces of an application on one or more of the cooperating computers of said cluster.

45. The memory medium of claim 37, wherein said first operation identifies said set at a first time based on said performance related metrics, and wherein step (d) changes said set at a second later time.

46. The memory medium of claim 37, wherein said fourth operation is performed only when a new value has been received or a request has been made to view the data.

47. The memory medium of claim 37, wherein said fourth operation forms said performance metrics as an aggregation of said values.

48. The memory medium of claim 47, wherein said fourth operation forms said performance metrics for each of said plurality of said cooperating computers.

49. The memory medium of claim 48, wherein said fourth operation is performed only when a new one of said values is received or a request to view the performance metric is received.

50. The memory medium of claim 48, wherein second operation receives said performance related values over a series of time intervals, and wherein said third

operation derives said performance metrics for n periods, of which the performance metrics of the n th period thereof includes an aggregate of the performance metrics for a current interval plus $n-1$ of the preceding intervals.

51. The memory medium of claim 50, wherein the performance metrics of the n th period of a preceding interval are discarded during a current interval.

52. The memory medium of claim 51, wherein said performance metrics for each of said periods include only a number and average of values received.

53. The memory medium of claim 37, wherein each of said performance metrics includes only a number and an average of values received.

54. The memory medium of claim 50, wherein said third operation forms said performance metrics as a data structure having n rows that contain the performance metrics of said n periods, respectively, wherein the performance metrics of the n th row of a preceding interval are discarded during a current interval, and wherein said n th row of the preceding interval is used as a first row in the current interval and the remaining ones of said n rows are shifted down one row position.